LEADU RESEARCH | HUMANING INTELLIGENCE ARCHITECTURE

AI AND JOBS

Why AI Overpromises and Underdelivers

Category errors, jagged HAICU, and the VOLTAGE fit required to convert intelligence into realized value

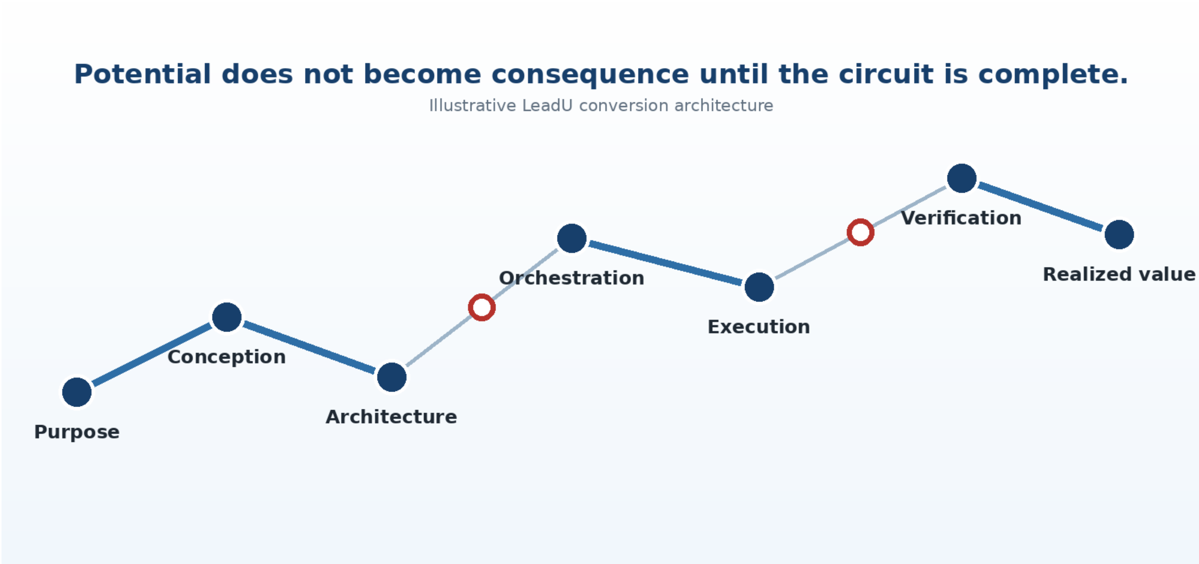

Cover figure. Illustrative LeadU conversion architecture. Red nodes represent consequential gaps that can interrupt realization.

"The debate is not only whether AI can do the work. It is whether human-AI configurations can complete the circuit from potential to consequence."

MIKE R. JAY, DEVELOPMENTALIST
Leadership University Online / LeadU

STATUS Working research and public orientation	**EDITION** v090826 | September 8, 2026
SCOPE AI, jobs, productivity, judgment formation, cognitive load, conversion inequality and shared prosperity	**LINEAGE** Supersedes v090726 | Human-skills evidence strengthening

What persists? • What changes? • What matters now?

Intelligence is necessary but never sufficient.

LeadU
LEADERSHIP UNIVERSITY ONLINE

Executive Summary

The AI-and-jobs debate is often organized around a false choice: either current employment stability proves that fears are overblown, or accelerating model capability proves that widespread displacement is imminent. Both positions can make the same modeling error. They take a true observation from one part of the system and use it to describe a future relationship that has not yet been established.

Current evidence is more specific. Stanford researchers report no widespread economy-wide displacement, yet the employment path for workers ages 22-25 in highly AI-exposed occupations is now 19 percent below that of less-exposed peers, with no comparable gap for experienced workers. Anthropic likewise finds no systematic rise in unemployment among highly exposed occupations, while reporting suggestive evidence of slower hiring for younger workers and a large gap between theoretical AI capability and observed use. [2, 3]

At the same time, AI use and scaling are not converting cleanly into enterprise value. McKinsey's full August 2026 survey of 1,719 participants across 97 nations reports that 80 percent experience individual productivity gains and 44 percent say AI is scaling across their enterprise, up from 38 percent a year earlier. Yet only 37 percent attribute a positive EBIT contribution to AI, essentially unchanged from 2025, and just 6 percent qualify as AI high performers. Those high performers are distinguished less by access to AI than by coherent workflow redesign, leadership behavior, measurement, risk management, human-in-the-loop design, and ownership of the change agenda. [1]

80%	37%	6%
report individual productivity gains	report a positive earnings contribution	qualify as AI high performers

McKinsey Global Survey, fielded May 4-June 8, 2026: 1,719 participants in 97 nations, respondent-reported and GDP-weighted. Survey evidence, not audited causal effects. [1]

This paper proposes that the missing conversion layer is not merely technical. Humans are jagged across knowledge, experience, judgment, timing, developmental complexity, relationships, and resources. AI is also jagged. Putting the two together does not automatically produce a complete capability circuit. In many cases, AI amplifies what the human-AI configuration already represents well while failing to recognize and bridge the very gaps for which the person most needs help.

> **LeadU working hypothesis**
>
> AI overpromises and underdelivers when available intelligence cannot be converted into realized consequence because the concurrent human-AI configuration cannot meet or regulate the VOLTAGE pressures presented at one or more consequential boundaries. External evidence establishes a conversion gap. VOLTAGE and jagged HAICU are LeadU candidate explanations for part of that gap, not externally validated causal findings.

Recent McKinsey research adds a further boundary. AI may raise the premium on judgment, contextual interpretation, coaching, exception handling, and quality control while absorbing the routine work through which novices historically developed those capabilities. Automation may also reduce task volume while increasing the cognitive intensity of the work that remains. Skill formation, sustainable load design, and human-agent performance therefore belong inside the conversion problem, not outside it. [13-17]

The policy and leadership issue is therefore larger than task performance. It concerns who can carry Purpose into realized value under CCR@VUCA, and who receives the benefit. This paper also follows organizational gains into labor and income consequences, ownership, purchasing power, and participation. That downstream loop is conditional: productivity does not automatically establish shared prosperity, and the sources do not establish inevitable unemployment or demand collapse.

Governing boundary: Do not collapse the chain
Exposure does not equal use. Use does not equal scaling. Scaling does not equal organizational absorption. Organizational absorption does not automatically equal realized value. Expected workforce change does not equal realized workforce change.

How to Read This Paper

This paper separates four kinds of claims. Observed evidence describes what a source measured. Forecasts describe what respondents or commentators expect. LeadU interpretation relates those findings through HIA, HAICU, CCR@VUCA, and VOLTAGE. Working hypotheses identify explanations that require further testing. Neutrality here does not mean giving equal weight to unequal evidence. It means preserving source scope, incentives, uncertainty, and boundaries before drawing a conclusion.

- Observed evidence: payroll records, model-use data, and organizational survey responses.
- Forecast evidence: employer expectations and scenario claims about what may happen next.
- LeadU interpretation: a governed attempt to relate separate components without treating any one source as sufficient.
- Working hypothesis: a candidate explanation that must remain open to correction by evidence and reality.

1. The Debate May Be Modeling the Wrong Problem

"We do not have to experience the cliff before we can Notice that the terrain has changed."

The most common category error begins with a defensible statement: "AI is not causing widespread job losses." The error occurs when that observation about now quietly becomes a prediction about next. Present employment data can constrain a forecast, but they cannot substitute for the causal bridge between current conditions and a future labor equilibrium.

The reverse error is equally possible. A forecast may begin with rapid increases in AI capability and assume that technical feasibility will immediately become deployment, organizational redesign, lower labor demand, and mass unemployment. That inference also skips the causal bridge. Adoption costs, regulation, trust, verification, institutional friction, human authority, and business-model change can slow, redirect, or block conversion.

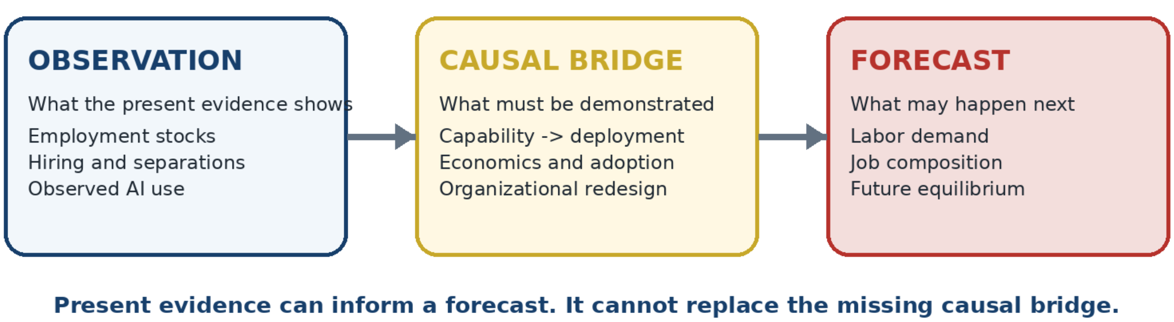

Figure 1. Present evidence can inform a forecast. It cannot replace the missing causal bridge. LeadU illustration.

Historical technological shifts remain useful evidence. They do not function as a law guaranteeing that this shift will follow the same sequence. AI is still technology, but its operating domain increasingly includes language, analysis, synthesis, coding, planning, judgment support, coordination, and portions of the learning processes through which people adapt to technological change. The historical analogy therefore has to be revalidated rather than simply repeated.

The sharper claim

This shift cannot be assumed to follow previous shifts. The reassuring evidence may be real. The reassurance projected from it may not be.

2. Employment Is a Stock; Labor Demand Moves Through Flows

A person does not have to be laid off for AI to be reducing labor demand. Employment is a stock shaped by flows into and out of jobs. The stock can remain stable while the flows underneath it change.

Stock-flow model

Future employment = current employment + hiring - separations

Conceptual identity. It does not specify the cause of any change in hiring or separations.

Incumbent workers can remain employed while firms hire fewer graduates, replace fewer departures, build smaller teams, remove developmental work, or demand more output from existing workers. A labor-market effect may therefore appear first as a narrowing entry ramp rather than a wave of layoffs.

The revised Stanford Digital Economy Lab paper illustrates this distinction. Its authors find no evidence of widespread economy-wide displacement. At the same time, employment among workers ages 22-25 in highly AI-exposed occupations is 19 percent below where it would be had it kept pace with less-exposed peers. The gap has widened since 2025, while experienced workers show no comparable pattern. [2]

Anthropic reports a compatible but separately measured signal: no systematic unemployment increase in highly exposed professions, but suggestive evidence that hiring of younger workers has slowed in exposed occupations. It also finds that actual AI coverage remains only a fraction of theoretical capability. [3]

A personal example in “People Have No Idea What’s About to Happen...” makes the non-hiring margin visible. The speaker reports using AI assistants for accounting-related work and says he has not needed to hire an accountant. This is a self-reported decision to forgo a service purchase, not a verified count of jobs eliminated or an audit of the assistants’ work. It illustrates why commissioned work and positions never opened matter alongside layoffs. [12]

Why the entry ramp matters

When developmental tasks disappear, the loss is not only a first job. The system may also lose the work through which beginners become experienced, organizations renew their talent, and tacit judgment is learned. A stable stock can conceal a weakening developmental pipeline.

The Judgment-Formation Paradox

McKinsey’s September 2026 collection argues that when AI removes rote tasks, the work that remains can require more judgment, critical thinking, and cognitive effort. Its July article on expertise makes the developmental problem explicit. Research, documentation, data cleanup, basic coding, and preliminary analysis are not only low-value tasks. They are among the experiences through which novices build instincts, context, pattern recognition, and judgment. [13, 15]

The article describes an “AI boost” for experienced developers who can steer and verify AI output and an “AI drag” for early-career developers who do not yet possess that judgment. The short-term incentive to hire senior people and automate junior work can therefore weaken the bottom of the talent pyramid on which future expertise depends. [15]

McKinsey's broader skills analysis estimates that current technologies could theoretically automate activities occupying about 57 percent of US work hours, while more than 70 percent of the skills employers seek appear in both automatable and nonautomatable work. The skills may endure, but where, how, and by how many people they are applied can still change. Technical potential remains different from adoption and from job loss. [14]

The answer is not to preserve busywork for its own sake. It is to redesign entry-level roles around deliberate practice, independent attempts, comparison with AI output, expert feedback, coaching, and increasing responsibility. In this design, AI can accelerate learning without removing the experiences through which judgment becomes reliable. [15]

> **Human skill premium boundary**
>
> Human skills at a premium does not establish human employment security. A skill can become more valuable per person while fewer people are employed to exercise it. The developmental ladder is not incidental inefficiency; it is part of the architecture through which competent judgment is produced.

3. AI Is Entering the Adaptation Loop

Earlier technologies changed tasks, industries, and economic structures. Humans adapted through education, apprenticeship, experimentation, coordination, and the creation of new roles. AI now operates on portions of those same cognitive and organizational processes. It can help people learn, plan, write, code, analyze, and coordinate, but it can also compress or remove the work through which those abilities were previously developed.

This does not mean human adaptation disappears. It means the adaptation mechanism itself is changing. A worker may need to learn a new domain while AI is simultaneously changing that domain, the learning materials, the employer's workflow, the number of available entry tasks, and the standard for acceptable output. Under CCR@VUCA, the requirement moves while the person is trying to meet it.

That creates two concurrent streams. AI can narrow performance differences within a bounded task by giving more people access to high-quality output. At the same time, it can widen differences in the ability to recognize what is missing, orchestrate multiple domains, verify consequences, capture economic value, and retain decision authority. Local capability compression can coexist with systemic conversion inequality.

> **Category boundary**
>
> Democratizing access to intelligence does not automatically democratize adaptability, agency, ownership, judgment, position, or economic benefit.

4. From Capability to Consequence

Neither current AI capability nor current employment is sufficient to forecast the future labor equilibrium. The consequential object is the conversion relationship among capability, economics, deployment, organizational design, competitive pressure, regulation, human adaptation, and verification. Within organizational design, this revision makes absorption and coherence explicit: the ability to coordinate workflow, authority, data, measurement, risk, roles, and change across the enterprise.

Illustrative labor-consequence model
AI capability x economics x adoption x organizational architecture (including absorption and coherence) x competitive pressure x regulation x human adaptation x verification -> labor consequence

This is an explanatory model, not a validated equation.

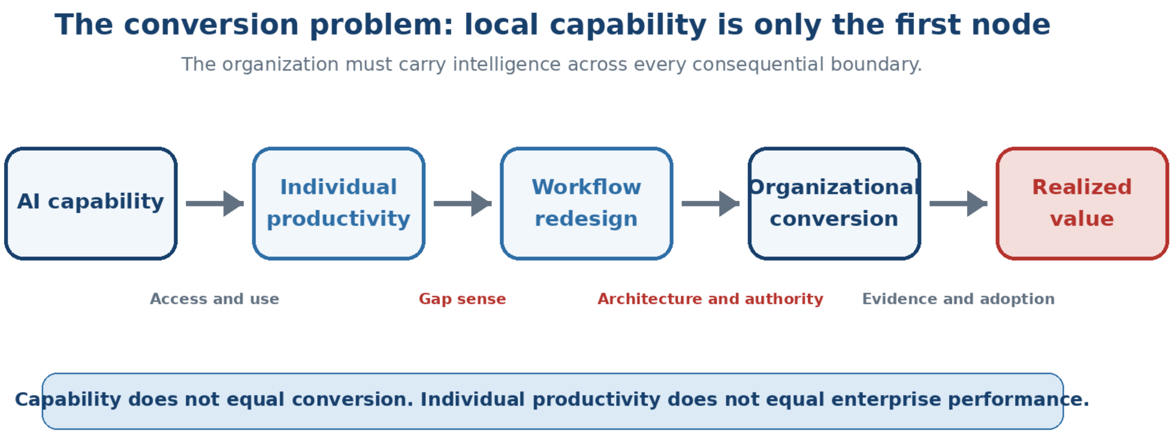

Figure 2. Local AI capability is only the first node. The organization must carry intelligence across every consequential boundary. LeadU illustration.

McKinsey’s full August 2026 report gives the conversion gap a stronger empirical surface. Forty-four percent of respondents now report enterprise-scale AI use, up from 38 percent a year earlier, while the share attributing a positive EBIT contribution to AI remains essentially unchanged at 37 percent. Just 6 percent qualify as high performers. Those high performers are 3.3 times more likely than others to intend fundamental business transformation, and nearly three-quarters report redesigning workflows around AI, compared with one-quarter of other respondents. [1]

McKinsey’s interpretation is organizational. High performers distinguish themselves by the coherence of their approach: workflow redesign, senior-leadership ownership, human-in-the-loop design, impact measurement, risk management, strategic workforce planning, transformation authority, and data or knowledge structures that reinforce one another. One in five respondents also reports that AI operating costs constrain use. These findings strengthen organizational absorption and coherence as a subcomponent of organizational architecture. [1]

Those findings do not independently validate LeadU’s VOLTAGE explanation. They do establish that access, local productivity, and even enterprise scaling are not sufficient. Something must carry the gain across the organization. That “something” includes workflow, authority, economics, evidence, readiness, relationships, measurement, cross-functional coherence, and the ability to Notice missing architecture.

The downstream extension in section nine asks what happens after organizational conversion succeeds. Labor and income consequences interact with ownership, prices, purchasing power, and demand, which can feed back into enterprise revenue and further AI deployment. This extension does not replace the labor-consequence model or assign a fixed direction to its effects.

Conversion boundary: AI use ≠ scaling ≠ organizational absorption ≠ realized value.

5. Organizations Were Built to Cross Jaggedness

Organizations exist, in part, because individual humans are jagged. No single person reliably carries all the knowledge, experience, judgment, relationships, resources, timing, and developmental complexity required to produce consequential outcomes alone. Accounting covers what engineering does not know. Engineering builds what sales cannot. Operations coordinates what strategy imagines. Governance protects what speed can overlook.

The organization's deeper business is not any one specialized task. It is the conversion of distributed potential into coordinated realized value. AI can now perform substantial work inside particular nodes, especially in technical fields. But performing a node is not the same as recognizing the missing nodes, sequencing them, managing their interdependence, preserving authority, and carrying the result into verified use.

> **The organizational boundary**
>
> AI can code. But code is not the organization's business. The organization's business is conversion.

Coding is an unusually favorable substrate for AI. It offers formal syntax, machine-readable state, testable outputs, repeatable environments, and comparatively explicit failure signals. Much of organizational life is less tractable. Purpose, stakeholder alignment, resource allocation, ambiguous requirements, tacit knowledge, cross-domain handoffs, Adoption, and accountability are not fully specified execution environments.

> **Author's field note**
>
> AI has already helped me finish work that the grind of detail had held back. I would not willingly give up that leverage. Yet the same work is repeatedly held at barriers where AI can perform a technical task but cannot reliably recognize, assemble, and verify the missing architecture without requiring me to learn enough of the unfamiliar domain to orchestrate it. Unrealized potential remains upstream because the circuit is incomplete.

McKinsey reports that 32 percent of respondents said their organizations declined to buy at least one software product or feature because agentic coding tools made internal development possible. This is strong evidence that coding is a favorable substrate and that AI is changing the build-versus-buy boundary. It is not evidence that software construction is equivalent to the organization's full conversion work. [1]
This produces a paradox that current AI marketing rarely admits: the user may need to know enough about the expertise they lack to successfully use AI to compensate for the expertise they lack. Expertise has not disappeared. It has moved upward into gap recognition, orchestration, evaluation, and consequence management.

6. Jagged HAICU and the VOLTAGE Boundary

HAICU means Human-AI Concurrent Understanding. It is not the simple addition of a human and a model. It is the evolving relational configuration through which the person and AI Notice, interpret, decide, act, verify, correct, and learn together while human authority over Purpose and consequential meaning remains protected.

Both sides of that configuration are jagged. A person may have high capability in Purpose, conception, synthesis, relationships, or developmental architecture and much less experience in software, finance, law, implementation, or distribution. A model may be strong in coding, recall, translation, or first-draft synthesis and weak in longitudinal continuity, local context, hidden dependencies, ambiguous consequence, verification, or cross-domain orchestration.

VOLTAGE @LeadU names the patterned developmental and systemic pressure created by differences across Vertical, Oblique, Lateral, Time, and AGE/Maturation complexities. It does not perform work. It conditions the potential for flow by shaping whether attention, energy, motivation, and coordinated behavior can move coherently toward RightACTION. [9]

- Vertical: the hierarchical complexity required by the work, especially across MHC L9-L16.
- Oblique: the integrative pathways across systems, often hidden as constraints or enablers.
- Lateral: the breadth of Knowledge, Skills, and Experience available to the situation.
- Time: sequencing, pacing, timing, and time-span alignment.
- AGE/Maturation: developmental readiness and experiential conditioning.

Canonical correction

Strictly in LeadU canon, neither a person nor an AI "has VOLTAGE" as though it were stored energy. VOLTAGE is the pressure field. This paper uses VOLTAGE fit as public shorthand for whether the concurrent human-AI configuration can meet and regulate that field well enough for flow to continue.

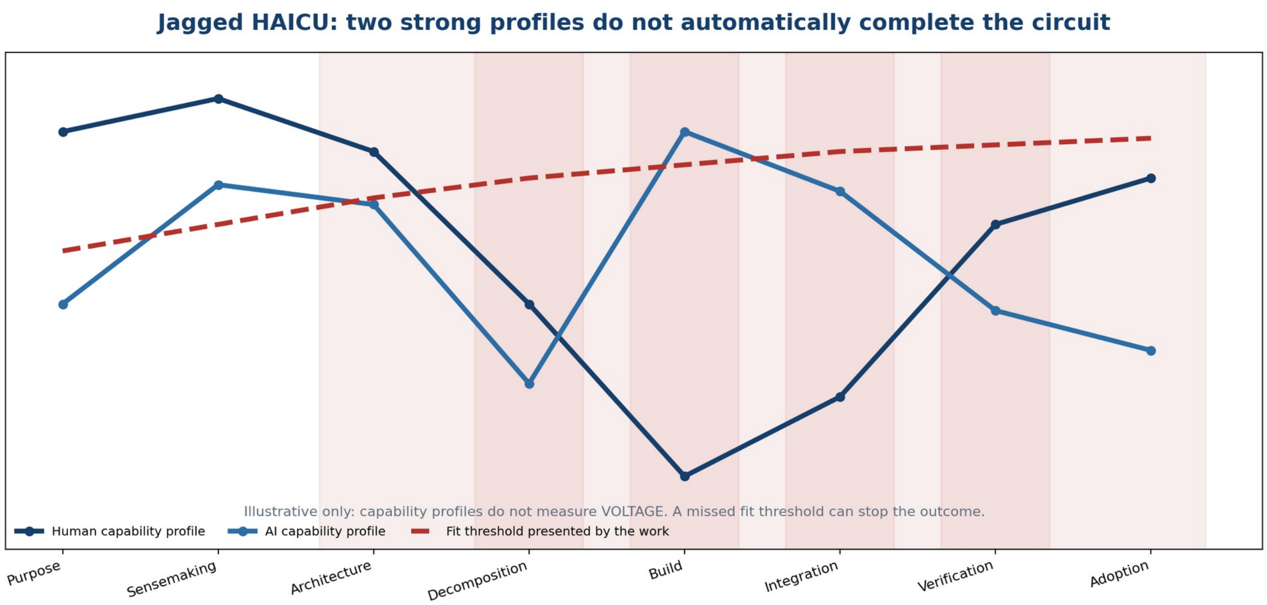

Figure 3. Illustrative jagged HAICU profile. The lines are conceptual, not measured. A single unbridged gap at a consequential node can interrupt the outcome.

The important unit is therefore not "human capability plus AI capability." It is whether the concurrent system can maintain outcome continuity across the next consequential boundary. A powerful model and a capable person can still form an insufficient circuit if their gaps overlap at architecture, decomposition, integration, verification, or Adoption.

LeadU candidate explanation

The productivity gap may partly be a VOLTAGE-fit gap: performance stalls not because intelligence is absent, but because the human-AI configuration cannot regulate the changing pressure field across the full path from Purpose to realized consequence.

The same report shows that local productivity can coexist with human strain. Forty-seven percent of midlevel managers and individual contributors report at least one AI-related negative effect, compared with 31 percent of executives and senior managers. Reported effects include stress, mental fatigue, output overload, pressure to take on more work, and concern about critical thinking. This does not validate VOLTAGE as the cause. It does show that benefit and friction can coexist inside the same human-AI configuration. [1]

McKinsey's work on brain-powered organizations extends the strain question from reported experience into work design. Automation can remove simpler duties while concentrating judgment, oversight, ambiguity, exception handling, interpersonal consequence, and accountability into each remaining hour. It may reduce total task volume while increasing cognitive intensity per hour. [16]

AI can also encourage cognitive offloading. If organizations do not preserve independent reasoning, deliberate practice, recovery, and sustained attention, the capabilities now commanding a premium can weaken rather than develop. [16]

In LeadU terms, the pressure field can intensify even when the task count falls. The remaining work may demand more Vertical complexity, more Oblique integration, broader Lateral knowledge, tighter Time judgment, and greater AGE/Maturation readiness. This is a LeadU interpretation through VOLTAGE. McKinsey does not measure VOLTAGE or establish it as the causal mechanism.

> **Load boundary**
>
> Task removal does not equal load reduction. Time saved does not equal cognitive capacity restored. Human work retained does not mean human work made sustainable.

7. Amplification Is Not Complementation

Current AI often amplifies capability more reliably than it complements jaggedness. A strong writer receives more writing leverage. A skilled analyst can perform more analysis. A software developer can use coding agents to build, test, and iterate faster. These are real gains. They are not yet the same as an AI that can recognize where the person lacks enough representation to formulate the missing work.

Technical practitioners frequently possess a "gap sense" that others do not. They can Notice that a project requires an API, repository, database, authentication layer, deployment pipeline, test harness, state model, or integration contract. Once named, AI may help build the missing structure. The nontechnical user may not know that the structure exists, much less how to specify or verify it.

Anthropic's January 2026 Economic Index supports part of this asymmetry. Claude use remained concentrated in coding and technical tasks, and model success declined as the human time and complexity of a task increased. The report also showed that nominal task coverage can overstate effective coverage when bottleneck tasks remain outside reliable AI performance. [4]

> **The democratization test**
>
> The user should not need the missing expertise in order to obtain the missing expertise. AI becomes broadly democratizing when it can increasingly recognize missing dependencies, explain them in fitted language, bridge what it can safely bridge, and identify where different human expertise or authority is still necessary.

That requirement must be governed. A gap-completing AI should not silently take authority over Purpose, meaning, risk acceptance, or consequential Adoption. The goal is not autonomous substitution for the person. It is complementary intelligence that strengthens human agency while making the missing architecture visible.

8. The Emerging Inequality Is a Conversion Inequality

The emerging divide may not be between people who have AI and people who do not. It may increasingly be between human-AI configurations that can complete capability circuits and those that cannot. Equal access to a frontier model can coexist with unequal gap sense, domain knowledge, time, money, organizational support, authority, confidence, infrastructure, and ability to absorb failure.

McKinsey's scaling data makes that inequality visible at the organizational level. Fifty-four percent of respondents from organizations with at least $1 billion in annual revenue report enterprise-scale AI, compared with about one-third of respondents from smaller organizations. Agent scaling rose to 40 percent among large organizations while remaining essentially flat at 22 percent among smaller organizations. The same models can therefore produce very different outcomes depending on the surrounding conversion architecture. [1]

PwC's 2026 AI Jobs Barometer reports widening divergence among firms. Companies in the most AI-exposed sectors showed higher productivity growth than less-exposed companies, but the top 20 percent of highly exposed firms produced a much larger "superstar" effect. The report also found a 62 percent average wage premium for AI skills. These are associations reported by an industry-linked source, not proof that AI alone caused the gains, but they are consistent with a conversion advantage accruing unevenly. [7]

The ILO's global exposure index shows a different inequality surface. One in four jobs has some potential GenAI exposure, but only 3.3 percent of global employment falls in the highest exposure category. Exposure is substantially higher in high-income economies than in low-income economies, and women are more represented in the highest-exposure group. The ILO's central conclusion is transformation rather than wholesale replacement, but transformation costs and benefits will not be evenly distributed. [6]

The World Economic Forum reports that 63 percent of surveyed employers identify skill gaps as a major barrier to business transformation. In its 100-worker illustration, 59 would need training by 2030 and 11 would be unlikely to receive it. Forty percent of employers anticipate workforce reductions where AI can automate tasks. These are employer expectations, but they show that the conversion problem is already being framed as a capability and organizational divide. [5]

The video "People Have No Idea What's About to Happen..." pushes this distinction toward productive ownership: using an AI service is not the same as holding claims on the income generated by an AI-enabled business. [12] A business may create productive assets while licensing its underlying models; it need not own the model itself. Access, income rights, operational control, and governance are separate questions. The proposed divide is therefore a risk to examine, not evidence that society must settle into two permanent classes.

> **Do not confuse local compression with systemic equality**
>
> AI can compress performance differences within a bounded task while widening differences in who can integrate, verify, own, govern, and economically capture the result.

9. Productive Abundance Does Not Establish Shared Prosperity

An organization can complete the circuit from AI capability to enterprise value while society fails to complete the circuit from enterprise value to broadly distributed human benefit. Jobs do more than organize production. They are also a route through which income, development, participation, and purchasing power reach people. When gains shift away from labor faster than prices, ownership, transfers, and new forms of participation adjust, productive potential may expand while the purchasing power of affected households weakens. This is the conditional risk examined here, not a forecast that automation must collapse demand.

From an individual decision to a system consequence

The business owner's accounting example contains both sides of the attraction. He reports relief from work that had burdened his business and recognizes that his decision also means another person was not hired. He then asks whether many individually rational substitutions could weaken the wage-supported demand on which businesses collectively rely. His account is a starting point for that question; it cannot establish the size of the effect beyond his business. [12]

The move from one business to an economy is precisely where another category error can occur. A cost reduction for one firm does not prove a loss of aggregate demand, just as that firm's improved margin does not prove a gain in shared prosperity. The result depends on what happens to the saving: whether it supports lower prices, new activity, investment, additional paid work, transfers, or income concentrated among owners. The next question is not simply whether the firm can produce with fewer people. It is whether the surrounding arrangements sustain participation and demand.

The demand-distribution feedback loop

The speaker describes a circle connecting wages, spending, business revenue, and renewed payments to workers. [12] This paper broadens that circle rather than treating wages as the only source of purchasing power. The proposed loop connects labor consequences to the distribution of income, purchasing power, demand, enterprise revenue, ownership and reinvestment, and further deployment. Existing ownership and institutional rules shape the relationships throughout; they are not merely an outcome at the end.

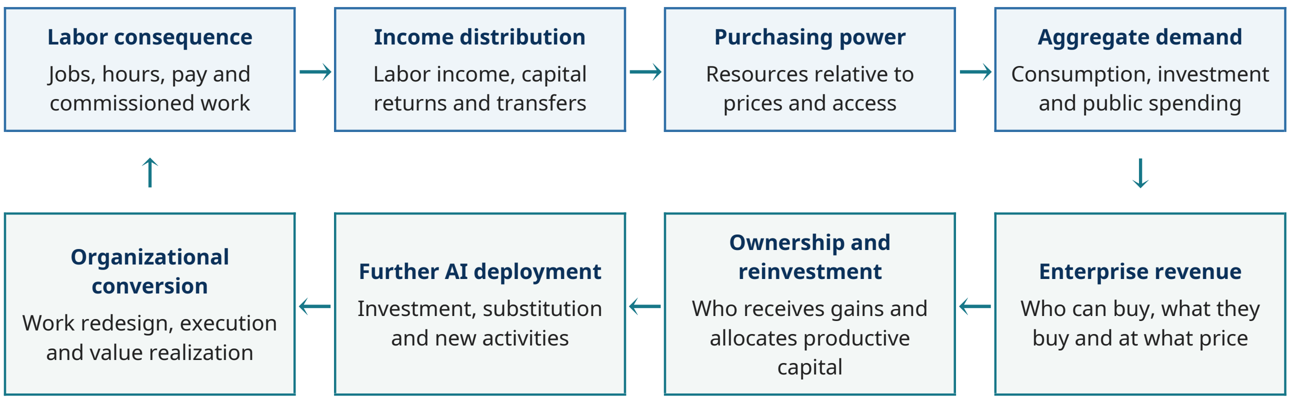

Figure 4. Conditional demand-distribution feedback. Arrows identify relationships to examine, not fixed causal effects. Ownership, prices, new work, transfers and public provision can alter the loop. LeadU illustration.

A firm may spend its savings elsewhere, expand output, create different work, or return income to owners. Households may receive capital income or transfers as well as pay. Investment and public spending also contribute to demand. These possibilities prevent the loop from becoming a story in which every job not created removes a customer permanently. They do not guarantee that gains reach the same people, places, or time horizons as the losses. The distribution and timing of adjustment remain the problem to investigate.

Cheap production does not settle access

The transcript distinguishes lower digital production costs from continuing physical constraints, particularly housing and land. Its point is that abundance does not automatically confer access when ownership and income remain unequal. [12] The stronger claim is conditional: lower prices help only insofar as people can obtain what they need on terms they can meet. Cheaper software alone does not establish secure housing, affordable infrastructure, or meaningful control over the systems on which a person depends.

Nor should digital abundance be treated as costless operation. McKinsey reports that AI-related operating costs constrain use for about one in five respondents' organizations. [1] That is an enterprise finding, not a household-access measure, but it cautions against assuming that an inexpensive additional output makes the surrounding service free. Production cost, market price, affordability, access rights, and reliability remain distinct.

Income support, ownership, and participation

The video imagines universal basic income as one response to extensive automation, while warning about "comfort without control." [12] This paper treats that as a scenario, not as the established purpose or inevitable result of income support. A payment may reduce hardship, sustain demand, and give people more room to choose. Whether it also increases or reduces dependence depends on its adequacy, security, conditions, and the wider rights and institutions around it.

Possible countervailing arrangements include new paid work, broader capital ownership, employee or community ownership, social dividends, public provision, and taxation or transfers. These are options for comparison, not findings from the video or endorsements of a single settlement. Their implementation, funding, accountability, and distributional effects would need separate examination. The pertinent questions include who receives the gains, who bears the transition costs, who can contest a decision, and who has a voice in the productive system.

The human question extends beyond the ability to buy. Income support may help people meet material needs without, by itself, establishing Purpose, contribution, belonging, ownership, or consequential authority. Conversely, paid employment is not the only possible source of those experiences. A Humaning approach should protect room for meaningful participation without requiring people to prove their worth by remaining economically indispensable to an employer.

What would support or weaken this warning?

The warning would gain support where reduced hiring, hours, or commissioned work coincides with deteriorating household purchasing power, concentrated returns, and inadequate alternative income or access. It would weaken where new work, lower essential costs, broader ownership, reliable transfers, or public provision offset those losses and strengthen participation. The transcript does not supply such population-level tests. McKinsey's organizational survey does not measure the household distribution of realized gains either. The extension identifies the relationships that need evidence rather than treating either source as proof of the complete loop. [1, 12]

> **A second conversion boundary**
>
> Productivity does not automatically establish prosperity. Cheap does not automatically mean accessible. Provision does not by itself establish freedom. Use does not confer ownership. Firm-level rationality does not establish system-level viability.

10. What Current Evidence Does and Does Not Establish

The sources below are not interchangeable. They observe different parts of the causal graph, use different methods, and carry different incentives. The table preserves their evidentiary contribution without allowing any one component to govern the whole interpretation.

Source	Evidence type	What it supports	What it does not establish
Stanford Digital Economy Lab, 2026	Administrative payroll evidence	No widespread economy-wide displacement; the employment path for workers ages 22-25 in highly exposed occupations is 19% below less-exposed peers; experienced workers show no comparable gap.	Does not establish that AI alone caused the full gap or that the future labor equilibrium is known. [2]
Anthropic, March 2026	Observed exposure framework and early labor evidence	Actual AI coverage remains far below theoretical capability; no systematic unemployment increase in highly exposed occupations; suggestive evidence of slower hiring for younger workers.	Provider-specific usage and early indicators do not establish economy-wide causal effects. [3]

Source	Evidence type	What it supports	What it does not establish
McKinsey, The state of AI in 2026, August 2026	Global organizational survey: 1,719 participants in 97 nations; respondent-reported and GDP-weighted	Eighty percent report individual productivity; 44 percent report enterprise-scale AI, up from 38 percent; 37 percent report positive EBIT, unchanged; 6 percent qualify as high performers. Larger organizations scale faster, 32 percent report build-versus-buy substitution through coding agents, and coherent redesign distinguishes high performers. Workforce expectations exceeded realized reductions. [1]	Does not establish audited causal effects, prove that AI caused EBIT or workforce changes, validate VOLTAGE as the conversion mechanism, or determine the future labor equilibrium.
Anthropic Economic Index, January 2026	Observed model-use patterns	Use remains concentrated in coding and other technical work. Success falls as task duration and complexity rise. Effective coverage is lower than task coverage when bottlenecks remain.	Claude usage is not representative of all AI, workers, or organizations. [4]
World Economic Forum, 2025	Employer forecast	Employers forecast 22% job churn by 2030, with 170 million roles created and 92 million displaced. Sixty-three percent identify skill gaps as a major transformation barrier.	This is an employer expectation survey, not an observed future fact. [5]
ILO and NASK, 2025	Global occupational exposure index	One in four jobs has some GenAI exposure; 3.3% fall in the highest exposure category; exposure is higher in high-income economies and for women. Transformation is more likely than full replacement.	Exposure does not equal use, automation, displacement, or harm. [6]
PwC, 2026	Industry-linked job and productivity analysis	AI-exposed companies show higher productivity growth, a 62% AI-skill wage premium, and a strong superstar effect among the highest-performing firms.	Associations do not prove AI causation or broad distribution of gains. [7]
Emad Mostaque interview, 2026	Bounded forecast stress case	Clarifies a plausible mechanism from capability to cost pressure, organizational redesign, and labor demand.	The two-year timetable is a forecast, not an established empirical finding. [8]
"People Have No Idea What's About to Happen..."	First-person account and economic commentary	Describes foregoing an accounting-service purchase and connects labor income, demand, ownership, and dependence as questions for examination. [12]	Does not establish net job loss, demand collapse, a timetable, or an inevitable income-support or ownership settlement.
McKinsey Global Institute, Agents, robots, and us, 2025	Technical-automation and skills analysis	Current technologies could theoretically automate activities occupying about 57% of US work hours. More than 70% of employer-sought skills occur in both automatable and nonautomatable work. [14]	Technical potential is not a forecast of job losses, adoption speed, or broad employment security.
McKinsey, Building expertise in the age of AI, 2026	Organizational synthesis with cited developmental studies	Routine entry work helps form judgment; experienced workers may gain an AI boost while novices face AI drag; attempt-then-check and coaching can preserve learning. [15]	Does not quantify economy-wide entry-level job loss or prove one developmental design works across all fields.

Source	Evidence type	What it supports	What it does not establish
McKinsey, Brain-powered organizations, 2026	Work-design synthesis with cited research	Automation can increase cognitive intensity per hour; role design, recovery, focus, independent thinking, and adaptive skills matter. [16]	Does not measure VOLTAGE or establish universal cognitive decline caused by AI.
McKinsey, Rewiring Talent to Value, 2026	Human-agent operating-model framework	Performance should be assessed at the coordinated human-agent system level, with clear roles, measures, accountability, feedback, escalation, and governance. [17]	Implementation guidance, not causal evidence that the framework guarantees superior outcomes.

11. A Bounded Forecast Stress Test

In an August 2026 interview with Peter McCormack, Emad Mostaque argued that the “economic life expectancy” of remote cognitive jobs could end within roughly two years as AI becomes able to replicate digital work at a fraction of human cost. His mechanism is competitive selection: organizations with fewer humans may become difficult for labor-intensive competitors to match. [8]

The timetable is not established by the cited evidence. Task exposure does not equal reliable end-to-end automation, and technical capability does not automatically become deployed organizational capability. The value of the argument is as a stress test. It asks whether current models of labor demand adequately represent the possibility that capability, cost, and competitive imitation could change organizational design faster than headline employment statistics reveal.

The reassuring camp and the catastrophic camp can therefore commit symmetrical errors. One extrapolates current labor-market stability while underweighting changing capability and competition. The other extrapolates capability while underweighting conversion friction, governance, trust, regulation, cost, and the jaggedness of real organizations.

> **Full boundary chain**
> Capability does not equal deployment. Deployment does not equal task automation. Task automation does not equal job elimination. Job elimination does not equal labor-market equilibrium. Expected workforce change does not equal realized workforce change.

McKinsey adds a useful calibration point. In the 2025 survey, 32 percent of respondents expected AI-driven workforce reductions during the following year. In 2026, only 14 percent reported actual reductions, while two-thirds reported little or no AI-related change. Yet 39 percent now expect reductions during the next year. This does not establish that future reductions will not occur. It establishes that expected workforce change is not the same as realized workforce change. [1]

12. What Must Change

Closing the gap requires more than larger models and more licenses. It requires design and governance that make conversion visible, attributable, verifiable, and fitted to actual people and situations. The same discipline must preserve how judgment is formed, manage the cognitive load of remaining work, and follow realized gains into their distribution rather than stopping at the enterprise boundary.

1 **Design AI for gap recognition, not only task completion.**
Models should identify likely missing dependencies, ask whether critical information or expertise is present, explain uncertainty, and show the path from a requested output to the surrounding architecture required for use.

2	**Make complementation a measured outcome.** Evaluation should ask whether AI bridges the user's jaggedness, not merely whether it improves a task score. Useful measures include missing-node detection, cross-domain continuity, handoff quality, verification success, and retained human authorship.
3	**Manage the human-agent system, not only its separate parts.** Define the agent's work, authority, performance measures, escalation conditions, evidence requirements, accountable human owner, correction process, and retirement or decommissioning criteria. Evaluate whether the complete workflow produces reliable, fitted, and human-authorized consequence. [17]
4	**Do not force every user to become a technical orchestrator.** AI literacy matters, but broad adoption will remain unequal if people must acquire substantial software, data, workflow, and integration expertise before AI can compensate for those same gaps.
5	**Redesign work without removing the developmental ladder.** Organizations should track which entry tasks, apprenticeships, and judgment-building experiences are being compressed. Replace incidental learning with protected practice, independent attempts, comparison with AI output, expert feedback, coaching, and progression toward independent judgment. Productivity gains that remove the path to experienced capability can create a delayed talent and governance failure. [15]
6	**Measure stocks, flows, scaling, absorption, and forecast error.** Monitor hiring, separations, replacement rates, entry-level openings, team size, task recomposition, workflow completion, enterprise scaling, organizational absorption, EBIT attribution, and expected versus realized workforce change. Headline unemployment and local productivity are insufficient.
7	**Preserve human authority over Purpose and consequence.** AI may surface options, relationships, risks, and missing architecture. It must not silently close inquiry into meaning, Purpose, acceptable consequence, or consequential Adoption.
8	**Build public and organizational conversion infrastructure.** Policy should support more than access to tools. People and smaller organizations need trusted technical translation, implementation support, verification, legal guidance, financing, and developmental help across the boundaries where potential currently stalls. The large-versus-small scaling gap should be monitored as an inequality indicator.
9	**Test the VOLTAGE-fit hypothesis without collapsing evidence.** Research should examine whether recurring failures cluster around Vertical, Oblique, Lateral, Time, and AGE/Maturation mismatches; whether task removal raises cognitive intensity per hour; whether organizational absorption and coherence mediate conversion; and whether improved HAICU design increases outcome continuity. [16]
10	**Follow the gains beyond the enterprise.** Track pay, hours, hiring, commissioned services, household income sources, essential costs, and the distribution of productive ownership alongside enterprise returns. Assess whether affected people gain real purchasing power and participation rather than assuming that firm-level savings reach them.
11	**Evaluate distribution arrangements and retained agency.** Compare new work, broader ownership, social dividends, transfers, and public provision for adequacy, durability, choice, and accountability. Income support should be assessed on its design and effects, not presumed either to deliver freedom automatically or to create dependence inevitably.

13. Conclusion: Completing Both Circuits

AI has not failed because it is unimpressive. It is impressive in ways that make the remaining boundary easier to miss. We see fluent outputs, fast code, rapid synthesis, and local productivity. We then assume that intelligence has reached the organization, the labor market, and the person's life as realized value.

Often it has not. The gain may stop at a missing dependency, an unseen technical boundary, a weak handoff, a timing mismatch, an absent verifier, a lack of authority, or a developmental gap. Current models often mirror and amplify the jagged configuration around them rather than complementing what is missing.

This is why AI can overpromise and underdeliver at the same time. The promise is expressed at the level of capability. Delivery is judged at the level of consequence. Between them lies the business of the organization, its ability to absorb and coherently coordinate change, the adaptation of the person, the economics of deployment, the governance of authority, and the VOLTAGE field through which MITEAM must flow.

> **The governing conclusion**
>
> The larger promise of AI is not intelligence availability. It is complementary human-AI capability that can carry Purpose through jagged boundaries into verified, fitted, and human-authorized consequence.

McKinsey's full report strengthens this boundary. Organizations are scaling AI faster than they are reporting financial returns, and the few high performers are distinguished by coherent redesign rather than tool accumulation. McKinsey describes the limiting factor as the organization's ability to absorb change. In LeadU terms, that is external convergence with the conversion problem, not independent proof that VOLTAGE fit is the measured cause. [1]

Organizational success is therefore necessary to the productive story, but it does not complete the distributional one. The further circuit concerns how gains reach people through income, ownership, access, participation, and the ability to influence the systems shaping their lives. That circuit can fail even when the first succeeds. It can also be strengthened through arrangements that share gains without making every person a technical operator or a business owner.

> **Four conversions, four separate tests**
>
> Capability → productive use
> Productive use → organizational value
> Organizational value → labor and income consequences
> Economic value → broadly shared human consequence

Human skills add a temporal and load constraint to both circuits. The work that becomes most valuable may be harder to cultivate if the learning tasks disappear, and the hours AI saves may return as denser judgment work rather than relief. A complete human-AI circuit must therefore preserve the development of judgment and the sustainability of those expected to exercise it. [15, 16]

The video supplies a personal illustration and a warning about ownership and dependence, not evidence that work must end or that an economy must run out of customers. [12] The proposed feedback loop keeps those questions conditional and makes the intervening arrangements visible. Neither a favorable productivity result nor an alarming forecast settles who benefits, who retains agency, or how the wider system adapts.

The lights do not come on because the generating plant is impressive. They come on when the productive circuit is complete. Yet completing that circuit does not settle who can afford the light, who owns the system, or who has a say in its operation. The next work is not merely to make AI more capable. It is to build Humaning architectures in which capability becomes RightACTION and the surrounding arrangements enable broadly shared benefit under real conditions.

LeadU Terms Used in This Paper

CCR@VUCA	Culture, Conditions, and Requirements operating within volatility, uncertainty, complexity, and ambiguity. It names the moving field in which awareness, readiness, and RightACTION are tested.
HAICU	Human-AI Concurrent Understanding. A longitudinal human-AI process and trust architecture for strengthening agency, preserving human authorship, and fitting intelligence to reality.
VOLTAGE	Vertical, Oblique, Lateral, Time, and AGE/Maturation complexities. VOLTAGE names patterned developmental and systemic pressure. It does not perform work; it conditions whether flow can move coherently toward RightACTION. [9]
VOLTAGE fit	Working public-language shorthand used in this paper for whether the human-AI configuration can meet and regulate the pressures presented across the five VOLTAGE dimensions. It is not a replacement definition for VOLTAGE.
Gap sense	The ability to Notice a missing dependency, boundary, capability, specialist, or architecture before the absent element has been explicitly represented.
HUMANING	The ongoing practice of living, relating, developing, and acting as a human under actual conditions rather than reducing the person to productivity, task output, or economic utility.
RightACTION	Action fitted to the person, situation, consequences, timing, readiness, and purpose. Speed or output alone does not establish RightACTION.

References

[1] **McKinsey & Company.** The state of AI in 2026: On the road to ROI. August 2026. Online

[2] Stanford Digital Economy Lab. Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence. Revised August 12, 2026. Online

[3] Anthropic. Labor market impacts of AI: A new measure and early evidence. March 5, 2026. Online

[4] Anthropic. Anthropic Economic Index report: Economic primitives. January 15, 2026. Online

[5] World Economic Forum. The Future of Jobs Report 2025. January 7, 2025. Online

[6] International Labour Organization and NASK. Generative AI and Jobs: A Refined Global Index of Occupational Exposure. May 20, 2025. Online

[7] PwC. AI reshapes global labour market into two distinct paths, rewarding human skills: PwC 2026 Global AI Jobs Barometer. 2026. Online

[8] The Peter McCormack Show. #206 - Emad Mostaque - “Your Job’s Economic Life Expectancy Ends In 2 Years”. August 28, 2026. Online

[9] Leadership University Online / LeadU. AGGI-POS Architectural Index | HIA/AAS Integration Edition, TPOV 72: VOLTAGE. v081826.

[10] Mike R. Jay, Developmentalist. Building an Evolving Humaning Intelligence Architecture. Working Research Draft v081726-R5.

[11] **McKinsey & Company.** Four perspectives on what matters now in AI. August 25, 2026. Online

[12] “People Have No Idea What’s About to Happen...” YouTube video. Online

[13] **McKinsey & Company.** The human skills at a premium in the age of AI. September 7, 2026. Online

[14] **McKinsey Global Institute.** Agents, robots, and us: Skill partnerships in the age of AI. November 25, 2025. Online

[15] **McKinsey & Company.** Building expertise in the age of AI: Who trains the next generation? July 14, 2026. Online

[16] **McKinsey & Company.** Five principles for designing brain-powered organizations. July 13, 2026. Online

[17] **McKinsey & Company.** Rewiring Talent to Value in the age of AI. June 18, 2026. Online

About LeadU

Leadership University Online develops Humaning Intelligence Architecture, HAICU, PA@LeadU, and related developmental systems for helping people and organizations work with intelligence without surrendering human authorship, Purpose, or consequential authority. HIA exists to govern human-AI concurrency so intelligence strengthens human agency and helps people carry RightACTION On Purpose in relation to reality.

Continue the inquiry at leadu.ai/news